

Wherefore, what is claimed is:

1. A computer-implemented process for generating a bi-level video, comprising using a computer to perform the following process actions:
 - 5 inputting digitized video comprising a sequence of video image frames;
 - converting the input video image frames into bi-level image frames;
 - and
 - 10 encoding the bi-level image frames.
2. The process of Claim 1, wherein the input video image frames have been compressed using a prescribed compression technique to facilitate storage or transmission, and wherein the process for generating the bi-level video further comprises an action of decompressing the input video image frames using a decompression technique appropriate for the prescribed compression technique, said decompression action being performed prior to the action of converting the input video image frames into bi-level image frames.
- 20 3. The process of Claim 1, wherein the inputted digitized video further comprises an audio component associated with each image frame, and wherein the process for generating a bi-level video further comprises the actions of:
 - 25 splitting the input video to separate the audio components from the image frames, prior to converting the image frames into bi-level image frames;
 - encoding the audio components using a low bit-rate audio compression technique;
 - 30 synchronizing and combining each audio component with the bi-level image frame that was derived from the input video image frame associated with the audio component.
4. The process of Claim 3, wherein the audio components have been compressed using a prescribed compression technique to facilitate storage or

transmission, and wherein the process for generating the bi-level video further comprises an action of decompressing the audio components using a decompression technique appropriate for the prescribed compression technique, said decompression action being performed prior to encoding the audio components using the low bit-rate audio compression technique.

5

10

5. The process of Claim 4, wherein the process actions of decompressing and encoding the audio components are performed concurrently with the process actions of converting the input video image frames into bi-level image frames and encoding the bi-level image frames.

15

6. The process of Claim 1, wherein the digitized video is characterized by pixels defined in terms either a color level or a gray scale level, and wherein the process action of converting the input video image frames into bi-level image frames, comprises the actions of:

20

ascertaining whether the input video image frames comprises pixels defined in terms of a color level; and
converting each input video image frame so as to define each pixel thereof in terms of one of a plurality of gray scale levels so as to produce a gray scale image frame whenever the pixels are defined in terms of color levels, wherein the particular gray scale level selected to define a pixel is based on the color level of that pixel.

25

7. The process of Claim 6, wherein the process action of converting the input video image frames into bi-level image frames, comprises, for each bi-level image frame, the actions of:

30

computing a threshold gray scale level;
assigning a first of two bi-level image frame colors to those pixels that exhibit a gray scale level exceeding said threshold by a prescribed upper-limit amount; and

assigning the second of said two colors to those pixel exhibiting a gray scale level that is less than said threshold by a prescribed lower-limit amount;

5 associating an indicator with those pixels that exhibit a gray scale level which does not exceed said threshold by the upper-limit amount and which are not less than said threshold by the lower-limit amount, wherein one of the two bi-level image frame colors is assigned to the pixels having said indicator associated therewith based on a predicted value established for the pixel as part of the performance of the process action for encoding the bi-level image frames.

10

8. The process of Claim 6, wherein the process action of converting the input video image frames into bi-level image frames, comprises, for each bi-level image frame, the actions of:

15 computing a threshold gray scale level;

assigning a first of two colors to those pixels that exhibit a gray scale level equal to or exceeding said threshold; and

assigning the second of said two colors to those pixel exhibiting a gray scale level that is less than said threshold.

20 9. The process of Claim 7, wherein the process action of computing a threshold gray scale level for a bi-level image frame, comprises the actions of:

(a) computing the mean of the gray scale levels of all the pixels of the image frame and designating the computed mean as a current preliminary gray scale level threshold value;

25 (b) computing the mean of all the gray scale levels of a first group of pixels having a gray scale level equal to or exceeding the currently-designated preliminary threshold value;

(c) computing the mean of all the gray scale levels of a second group of pixels having a gray scale level less than the currently-designated 30 preliminary threshold value;

(d) computing the average of the mean gray scale levels computed for the first and second groups of pixels;

(e) designating the computed average gray scale level of the two pixel groups as the current preliminary gray scale level threshold value in lieu of the previously-designated value;

5 (f) repeating actions (b) through (e), until the computed average gray scale level does not change; and

(g) designating the last-computed average gray scale level as the threshold gray scale level for the bi-level image frame under consideration.

10

10. The process of Claim 7, wherein the process action of computing a threshold gray scale level for a bi-level image frame, comprises the actions of:

15 (a) computing the mean of the gray scale levels of all the pixels of the image frame and designating the computed mean as a current preliminary gray scale level threshold value;

(b) computing the mean of all the gray scale levels of a first group of pixels having a gray scale level equal to or exceeding the currently-designated preliminary threshold value;

20 (c) computing the mean of all the gray scale levels of a second group of pixels having a gray scale level less than the currently-designated preliminary threshold value;

(d) computing the average of the mean gray scale levels computed for the first and second groups of pixels;

(e) designating the computed average gray scale level of the two pixel groups as the current preliminary gray scale level threshold value in lieu of the previously-designated value;

25 (f) repeating actions (b) through (e), until the computed average gray scale level does not change;

(g) inputting a user-supplied gray scale level threshold

30 adjustment value;

5 (h) computing the sum of said threshold adjustment value and the last-computed average gray scale level; and

5 (i) designating the sum of the threshold adjustment value and the last-computed average gray scale level as the threshold gray scale level for the bi-level image frame under consideration.

10 11. The process of Claim 7, wherein the first of the two bi-level image frame colors is lighter than the second of the colors.

10 12. The process of Claim 7, wherein the first of the two bi-level frame colors in white, and the second of said colors is black.

15 13. The process of Claim 7, wherein the upper-limit amount equals the lower-limit amount.

15 14. The process of Claim 7, wherein the upper-limit amount varies within a range of 0 to about 5 gray scale levels.

20 15. The process of Claim 7, wherein the lower-limit amount varies within a range of 0 to about 5 gray scale levels.

25 16. The process of Claim 7, further comprising a process action, which is performed prior to assigning any bi-level image frame color to the pixels of a frame of the input video, of reducing a flickering effect in the bi-level video caused when correspondingly-located pixels in a series of consecutive frames which depict the same unchanged portion of a captured scene are varyingly assigned one or the other of two bi-level image frame colors owing to a change in the lighting conditions between the times the individual frames of the series were captured.

17. The process of Claim 16, wherein the process action of reducing a flickering effect, comprises, for each pixel in each consecutive frame of the input video starting with the second frame, the actions of:

5 computing a Laplacian of the pixel under consideration in a current frame and of each pixels contained in the current frame within an region centered on the pixel under consideration;

10 computing Laplacians for each correspondingly-located pixel in the frame immediately preceding the current frame in the frame sequence which are contained within a region corresponding in location to said region in the current frame;

15 computing the difference between the Laplacian computed for each pair of correspondingly-located pixels in the current frame and its immediately preceding frame;

20 summing the absolute value of the computed differences, and designating the sum as a SAD of the pixel under consideration;

25 determining whether the SAD of the pixel under consideration is greater than a prescribed dissimilarity threshold; and

whenever the SAD is equal to or less than a prescribed dissimilarity threshold,

30 assigning to the pixel under consideration an indicator indicating the pixel is similar to the correspondingly-located pixel in the immediately preceding frame,

computing the difference between the gray scale level threshold computed for the current frame and the gray scale level threshold computed for immediately preceding frame,

adding the computed difference to the gray level value assigned to the pixel corresponding in location to the pixel under consideration in the immediately preceding frame to create a compensated gray level value, and

assigning the compensated gray level value to the pixel under consideration in lieu of the value previously assigned to the pixel.

18. The process of Claim 17, wherein said region centered on the pixel under consideration is square and contains 81 pixels.

5 19. The process of Claim 17, wherein said prescribed dissimilarity threshold is set between approximately 1.0 and 3.0.

10 20. The process of Claim 17, further comprises a process action of assigning to the pixel under consideration an indicator indicating this pixel is dissimilar to the correspondingly-located pixel in the immediately preceding frame whenever the SAD is greater than the dissimilarity threshold.

15 21. The process of Claim 20, wherein the input video frame under consideration depicts the upper body of a person, and wherein the process of generating a bi-level video further comprises the actions of, for each pixel row in a frame under consideration:

20 identifying the two outermost pixels that are assigned an indicator indicating the pixel is dissimilar to the correspondingly-located pixel in the immediately preceding frame, for those pixel rows having at least two pixel marked as being dissimilar;

25 identifying the intervening pixels that are assigned an indicator indicating the pixel is similar to the correspondingly-located pixel in the immediately preceding frame and which are between the identified outermost pixels assigned an indicator indicating the pixel is dissimilar to the correspondingly-located pixel in the immediately preceding frame;

30 assigning to each identified intervening pixel an indicator indicating this pixel is dissimilar to the correspondingly-located pixel in the immediately preceding frame; and

reassigning to each identified intervening pixel its original gray scale value.

30

22. The process of Claim 20, further comprises the process actions of:

identifying for the pixel under consideration whether it is assigned an indicator indicating it is similar to the correspondingly-located pixel in the immediately preceding frame;

5 whenever the pixel under consideration is identified as having been assigned an indicator indicating it is similar to the correspondingly-located pixel in the immediately preceding frame, ascertaining whether less than a prescribed number of corresponding-located pixels in frames sequentially preceding the frame under consideration are each marked as dissimilar; and

10 whenever less than the prescribed number of corresponding-located pixels in frames sequentially preceding the frame under consideration are each marked as dissimilar,

15 assigning an indicator indicating the pixel under consideration is dissimilar to the correspondingly-located pixel in the immediately preceding frame in lieu of an indicator indicating similarity, and

15 reassigning the original gray scale value associated with the pixel under consideration.

23. The process of Claim 20, further comprises the process actions of:

20 identifying in the frame under consideration each region of substantially contiguous pixels that has been assigned an indicator indicating they are dissimilar to their correspondingly-located pixels in the immediately preceding frame; and

25 for each pixel located immediately adjacent each identified dissimilarity region,

assigning an indicator indicating the pixel is dissimilar to the correspondingly-located pixel in the immediately preceding frame in lieu of an indicator indicating similarity, and

reassigning the original gray scale value associated with the pixel.

24. The process of Claim 1, wherein the process action of encoding the bi-level image frames, comprises an action of encoding the bi-level frames employing an adaptive context-based arithmetic coding technique.

5 25. The process of Claim 1, wherein the process action of encoding the bi-level image frames, comprises the actions of, for each bi-level image frame:

encoding the bi-level frame employing an adaptive context-based arithmetic coding technique which utilizes a $\frac{1}{2}$ probability table;

10 encoding the bi-level frame employing an adaptive context-based arithmetic coding technique which utilizes a pre-trained probability table;

determining whether the encoding utilizing a $\frac{1}{2}$ probability table or the encoding utilizing a pre-trained probability table produces a bit stream with the lowest bit rate;

15 designating the bit stream produced using the probability table determined to produce the lowest bit rate as representing the encoded bi-level frame; and

20 incorporating in the bit stream designated as representing the encoded bi-level frame an indicator of the type of probability table used to encode the bit stream.

26. The process of Claim 7, wherein the process action of encoding the bi-level image frames, comprises an action of encoding the bi-level frames employing an adaptive context-based arithmetic coding technique that comprises predicting whether a pixel under consideration exhibits one or the other of the

25 two bi-level image frame colors based on the values of a prescribed pattern of previously-predicted pixels and then compares the predicted pixel value to the actual value of the pixel and encodes the bi-level image frame under

30 consideration by indicating those pixels for which the predicted pixel value does not match the actual pixel value, with the exception of those pixels having said indicator associated therewith that indicates the pixel exhibits a gray scale level which does not exceed said threshold by the prescribed upper-limit amount and

which are not less than said threshold by the prescribed lower-limit amount, said excepted pixels being presumed to have a predicted value that matches its actual value.

5 27. The process of Claim 20, wherein the process action of encoding the bi-level image frames, comprises the actions of:

designating the first bi-level image frame and frames in the frame sequence occurring at a prescribed interval as I-frames;

10 designating the bi-level image frames in the frame sequence falling between a pair of consecutive I-frames, as p-frames;

determining the smallest bounding box that will surround all regions of dissimilarity in each p-frame;

15 encoding each I-frame using an adaptive context-based arithmetic coding technique; and

encoding each p-frame using an adaptive context-based arithmetic coding technique.

20 28. The process of Claim 27, wherein the process action of encoding each I-frame, comprises the actions of:

encoding the I-frame employing an adaptive context-based arithmetic coding technique which utilizes a $\frac{1}{2}$ probability table;

encoding the I-frame employing an adaptive context-based arithmetic coding technique which utilizes a pre-trained probability table;

25 determining whether the encoding utilizing a $\frac{1}{2}$ probability table or the encoding utilizing a pre-trained probability table produces a bit stream with the lowest bit rate;

designating the bit stream produced using the probability table determined to produce the lowest bit rate as representing the encoded I-frame; and

incorporating in the bit stream designated as representing the encoded I-frame an indicator of the type of probability table used to encode the bit stream.

5 29. The process of Claim 27, wherein the process action of encoding each p-frame, comprises the actions of:

 encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a $\frac{1}{2}$ probability table;

10 encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a pre-trained probability table;

15 determining whether the encoding utilizing a $\frac{1}{2}$ probability table or the encoding utilizing a pre-trained probability table produces a bit stream with the lowest bit rate;

 designating the bit stream produced using the probability table determined to produce the lowest bit rate as representing the encoded p-frame; and

20 incorporating in the bit stream designated as representing the encoded p-frame an indicator of the type of probability table used to encode the bit stream and the size and location of the bounding box.

30 30. The process of Claim 28, wherein the process action of encoding each p-frame, comprises the actions of:

 encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a $\frac{1}{2}$ probability table;

 encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a first pre-trained probability table;

encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a second pre-trained probability table which is identical to that employed to encode the last preceding I-frame in relation to the p-frame under consideration;

5

determining whether the encoding utilizing the $\frac{1}{2}$ probability table, the first pre-trained probability table, or the second pre-trained probability table produces a bit stream with the lowest bit rate;

10

designating the bit stream produced using the probability table determined to produce the lowest bit rate as representing the encoded p-frame; and

incorporating in the bit stream designated as representing the encoded p-frame an indicator of the type of probability table used to encode the bit stream and the size and location of the bounding box.

15

31. The process of Claim 28, wherein the process action of encoding each p-frame, comprises the actions of:

encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a $\frac{1}{2}$ probability table;

20

encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a first pre-trained probability table;

25

encoding the portion of the p-frame contained within said bounding box employing an adaptive context-based arithmetic coding technique which utilizes a second pre-trained probability table which is identical to the final probability table resulting from the complete encoding of the last preceding I-frame in relation to the p-frame under consideration;

30

determining whether the encoding utilizing the $\frac{1}{2}$ probability table, the first pre-trained probability table, or the second pre-trained probability table produces a bit stream with the lowest bit rate;

designating the bit stream produced using the probability table determined to produce the lowest bit rate as representing the encoded p-frame; and

incorporating in the bit stream designated as representing the encoded p-frame an indicator of the type of probability table used to encode the bit stream and the size and location of the bounding box.

32. The process of Claim 17, further comprising a process action of controlling the bit rate at which the bi-level video image frames are encoded so as to maintain the average bit rate to less than or approximately equal to a target bit rate consistent with low bandwidth devices.

33. The process of Claim 32, wherein the process action of controlling the bit rate at which the bi-level video image frames are encoded, comprises the actions of:

establishing a rate control table having a hierarchical sequence of rate control scale factors each of which is assigned values for at least one of (i) said prescribed upper-limit and lower-limit amounts, and (ii) said prescribed dissimilarity threshold, wherein the values assigned to a rate control scale factor higher in the hierarchy of the sequence are the same or larger than the values assigned to a lower level scale factor;

initially selecting one of the rate control scale factors and then utilizing any values assigned thereto for said prescribed upper-limit and lower-limit amounts and said prescribed dissimilarity threshold in converting image frames of the input video into bi-level image frames;

establishing a buffer of a prescribed size for storing the bits produced as the bi-level video image frames are encoded, until their transfer for further processing;

determining whether the number of bits stored at any time in the buffer exceed the half-size of the buffer by more than a first prescribed amount

or are less than the half-size of the buffer by more than a second prescribed amount;

whenever the number of bits stored at any time in the buffer exceed the half-size of the buffer by more than the first prescribed amount,
5 selecting the rate control scale factor one level up from the previously selected level and employing any values assigned to the newly-selected scale factor for the upper-limit and lower-limit amounts and the dissimilarity threshold in converting the next, previously-unconverted, image frame of the input video into a bi-level image frame, unless the top scale factor level is the currently-selected level in which case the currently-employed values of the upper-limit and lower-limit amounts and the dissimilarity threshold are used; and

whenever the number of bits stored at any time in the buffer are less than the half-size of the buffer by more than the second prescribed amount, selecting the rate control scale factor one level down from the previously selected level and employing any values assigned to the newly-selected scale factor for the upper-limit and lower-limit amounts and the dissimilarity threshold in converting the next, previously-unconverted, image frame of the input video into a bi-level image frame unless the lowest scale factor level is the currently-selected level in which case the currently-employed values of the upper-limit and lower-limit amounts and the dissimilarity threshold are used.

34. The process of Claim 33, further comprising the process actions of: determining whether the number of bits stored at any time in the buffer exceed the size of the buffer; and

25 whenever the number of bits stored at any time in the buffer exceed the size of the buffer,

deleting the bits associated with the last-converted image frame from the buffer,

30 selecting the top rate control scale factor level, and employing any values assigned to the newly-selected scale factor for the upper-limit and lower-limit amounts and the dissimilarity threshold

in converting the next, previously-unconverted, image frame of the input video into a bi-level image frame, unless the top scale factor level is the currently-selected level in which case the currently-employed values of the upper-limit and lower-limit amounts and the dissimilarity threshold are used.

5

35. The process of Claim 1, wherein the bi-level video is to be transmitted over a medium requiring a prescribed transmission encoding scheme, and wherein the process for generating a bi-level video further comprises a process action of packetizing the bi-level image frames in accordance with the prescribed packetizing scheme applicable to the medium over which the bi-level video is to be transmitted.

10

36. The process of Claim 4, wherein the bi-level video is to be transmitted over a medium requiring a prescribed transmission packetizing scheme, and wherein the process for generating a bi-level video further comprises a process action of packetizing the combined bi-level frames and audio components in accordance with the prescribed packetizing scheme applicable to the medium over which the bi-level video is to be transmitted.

15

37. A system for generating a bi-level video, comprising:
a video camera;
a sending device capable of transmitting encoded bi-level image frames and audio data and having a general purpose computing device; and
a computer program comprising program modules executable by the computing device, wherein the computing device is directed by the program modules of the computer program to,

input digitized video captured using said video camera, said video comprising a sequence of video image frames each with an associated audio component;

20

split the input video to separate the audio components from the image frames;

convert the input video image frames into bi-level image frames;
5 encode the bi-level image frames;
encode the audio components;
synchronize and combine each audio component with the bi-level image frame that was derived from the input video image frame associated with the audio component to produce a sequence of encoded bi-level video frames;
transmit each encoded bi-level video frame in sequence.

10

38. The system of Claim 37, wherein the video camera is a digital video camera and the digitized video is input into the computing device directly from said camera through an appropriate interface.

15

39. The system of Claim 37, wherein the video camera is a non-digital video camera and the images captured by the camera are digitized prior to being input into the computing device.

20

40. The system of Claim 37, further comprising a receiving device which is capable of receiving and decoding bi-level video frames, displaying bi-level video image frames and playing audio data.

25

41. The system of Claim 37, wherein the sending device is one of a (i) hand-held Personal Computer (PC), (ii) palm-sized PC, or (iii) mobile telephone.

42. A computer-readable medium having computer-executable instructions for generating a bi-level video, said computer-executable instructions comprising:

30

inputting digitized video comprising a sequence of video image frames;

converting the input video image frames into bi-level image frames;
and

compressing the bi-level image frames using an adaptive context-based arithmetic coding technique.

5

43. The computer-readable medium of Claim 42, further comprising an instruction for decompressing each input video image, prior to converting each image into a bi-level image, whenever the input video image frames are received in a compressed condition.

10

44. A computer-implemented process for playing a bi-level video, comprising using a computer to perform the following process actions:

receiving the bi-level video in the form of a sequence of frames each of which comprises a bi-level image;

15

decoding each bi-level image frame;

rendering each decoded bi-level image frame for display by assigning a first of two prescribed colors to pixel exhibiting a first of two binary values and the second of said prescribed colors to all remaining pixels of the image frame; and

20

displaying each rendered bi-level image frame using the assigned colors and in the sequence in which it was received.

45. The process of Claim 44, wherein the first prescribed color is white and the second prescribed color is black.

25

46. The process of Claim 44, wherein the second prescribed color is darker than the first prescribed color.

30

47. The process of Claim 44, wherein the second prescribed color contrasts the first prescribed color to the extent that it is readily discernable by a person viewing the bi-level video.

48. The process of Claim 44, wherein the bi-level video is transmitted over a medium requiring a prescribed transmission packetizing scheme, and wherein the process for playing a bi-level video further comprises a process action of unpacking the received frames of the b-level video using a unpacking procedure appropriate for the prescribed transmission encoding scheme.

49. The process of Claim 44, wherein the bi-level video image frames have been encoded using an adaptive context-based arithmetic coding technique, and wherein the process action of decoding each bi-level image frame comprises using an adaptive context-based arithmetic decoding technique.

50. The process of Claim 49, wherein the adaptive context-based arithmetic decoding technique comprises using a particular initial probability table for each bi-level image frame to predict whether each pixel in the image frame under consideration exhibits one or the other of the two bi-level image frame colors based on the values of a prescribed pattern of previously-predicted pixels, and wherein the process action of decoding the bi-level image frames, comprises the actions of:

10 receiving in conjunction with each encoded bi-level image frame a indicator specifying of the type of initial probability table that was used to encode that image frame; and

15 employing the indicated type of probability table as the initial probability table in the decoding of the bi-level image frame under consideration.

20 51. The process of Claim 44, wherein the bi-level video image frames have been encoded using an adaptive context-based arithmetic coding technique, and wherein the process action of decoding each bi-level image frame comprises the actions of:

designating the first bi-level image frame received and frames in the frame sequence occurring at a known prescribed interval as I-frames;
designating the bi-level image frames in the frame sequence falling between a pair of consecutive I-frames, as p-frames;
5 decoding each I-frame using an adaptive context-based arithmetic decoding technique; and
 decoding each p-frame using an adaptive context-based arithmetic decoding technique.

10 52. The process of Claim 51, wherein the adaptive context-based arithmetic coding technique used to encode the bi-level image p-frames comprises determining the smallest bounding box that will surround all regions of dissimilarity in each p-frame, encoding just the portion of the p-frame contained within said bounding box, and incorporating in the bit stream designated as 15 representing the encoded p-frame an indicator specifying the size and location of the bounding box in relation to an overall bi-level image frame, and wherein the process action of decoding each p-frame, comprises the actions of:

20 receiving in conjunction with each encoded p-frame the indicator specifying of the size and location of the bounding box in relation to the overall bi-level image frame; and

 assuming all pixels located outside the bounding box in the overall image frame are have the same value as those in the corresponding locations in the last preceding frame.

25 53. The process of Claim 44, wherein each bi-level video frame further comprises an audio data component, and wherein the process for playing a bi-level video further comprises the process actions of:

 splitting each bi-level video frame into its constituent audio data component and bi-level image frame, prior to processing the image frame; and

playing the audio data associated with a bi-level video frame concurrently with performing the process action of displaying the bi-level image frame associated with that bi-level video frame.

5 54. The process of Claim 53, wherein the audio data component of each bi-level video frame has been compressed using a audio compression technique, and wherein the process for playing a bi-level video further comprises a process action of decompressing the audio data associated with each frame of the bi-level video, prior to playing the audio data, using an audio decompression technique appropriate for the audio compression technique employed to compress the audio data component.

10 55. A system for playing a bi-level video, comprising:
15 a receiving device having a display capable of displaying bi-level images and playing audio data and having a general purpose computing device; and
20 a computer program comprising program modules executable by the computing device, wherein the computing device is directed by the program modules of the computer program to,
25 receive the bi-level video in the form of a sequence of frames each of which comprises a bi-level image each bi-level video frame and an audio data component,
 split each bi-level video frame into its constituent audio data component and bi-level image frame,
 decode each bi-level image frame,
 render each decoded bi-level image frame for display by assigning a first of two prescribed colors to pixel exhibiting a first of two binary values and the second of said prescribed colors to all remaining pixels of the image frame,
30 display each rendered bi-level image frame using the assigned colors and in the sequence in which it was received, and

play the audio data associated with each bi-level video frame concurrently with displaying the bi-level image frame associated with that bi-level video frame.

5 56. The system of Claim 55, wherein the audio data component of
each bi-level video frame has been compressed using an audio compression
technique, and wherein the computer program further comprises a program
module for decompressing the audio data associated with each frame of the bi-
level video, prior to playing the audio data, using an audio decompression
technique appropriate for the audio compression technique employed to
10 compress the audio data component.

57. The system of Claim 55, further comprising a sending device which is capable of sending audio data.

58. The system of Claim 55, wherein the receiving device is one of a (i) hand-held Personal Computer (PC), (ii) palm-sized PC, or (iii) mobile telephone.

59. A computer-readable medium having computer-executable instructions for displaying a bi-level video, said computer-executable instructions comprising:

receiving the bi-level video in the form of a sequence of frames each of which comprises a bi-level image encoded using an adaptive context-based arithmetic coding technique;

decoding each bi-level image frame using an adaptive context-based arithmetic decoding technique;

rendering each decoded bi-level image frame for display by assigning a first of two prescribed colors to pixel exhibiting a first of two binary values and the second of said prescribed colors to all remaining pixels of the image frame; and

displaying each rendered bi-level image frame using the assigned colors and in the sequence in which it was received.